Experimental Bayesian Quantum Phase Estimation on a Silicon Photonic Chip

S. Paesani,[1,*] A.A. Gentile,[1,*] R. Santagati,[1] J. Wang,[1] N. Wiebe,[2,†] D.P. Tew,[3] J.L. O'Brien,[1] and M.G. Thompson[1,‡]

[1] *Quantum Engineering Technology Labs, H. H. Wills Physics Laboratory and Department of Electrical and Electronic Engineering, University of Bristol, BS8 1FD, UK.*
[2] *Quantum Architectures and Computation Group, Microsoft Research, Redmond, Washington 98052, USA*
[3] *School of Chemistry, University of Bristol, Bristol BSB 1TS, UK*

(Dated: March 16, 2017)

Quantum phase estimation is a fundamental subroutine in many quantum algorithms, including Shor's factorization algorithm and quantum simulation. However, so far results have cast doubt on its practicability for near-term, non-fault tolerant, quantum devices. Here we report experimental results demonstrating that this intuition need not be true. We implement a recently proposed adaptive Bayesian approach to quantum phase estimation and use it to simulate molecular energies on a Silicon quantum photonic device. The approach is verified to be well suited for pre-threshold quantum processors by investigating its superior robustness to noise and decoherence compared to the iterative phase estimation algorithm. This shows a promising route to unlock the power of quantum phase estimation much sooner than previously believed.

arXiv:1703.05169v1 [quant-ph] 15 Mar 2017

Introduction—Quantum algorithms promise exponential advantages over their classical counterparts, allowing the possibility to accomplish tasks otherwise unachievable on a classical computer [1, 2]. A fundamental tool in quantum computing is the quantum phase estimation algorithm (PEA), necessary for harnessing many of its main applications, e.g factorization of large numbers [1, 3–6] and simulation of molecular properties [7–11]. An efficient implementation of PEA will thus be a crucial subroutine for quantum computers. Kitaev's iterative phase estimation algorithm (IPEA) [12] and its adaptive version [13, 14] have been employed in proof-of-principle implementations of PEA, as they solely rely on a relatively small number of qubits and logic gates [4–10]. However, they require exponentially long coherence of the quantum hardware and are very susceptible to experimental noise [9, 14–16]. This means that conventional quantum phase estimation algorithms rapidly become impractical if the quantum computer is not fully error-corrected, limiting the feasibility in near-term, pre-fault-tolerant quantum machines.

A new efficient Bayesian phase estimation algorithm, called *rejection filtering phase estimation* (RFPE), has been recently proposed to overcome this issue [16]. The algorithm applies an approximate form of Bayesian inference to efficiently estimate the correct eigenphase. Theoretical results suggest that RFPE has a number of desirable features: high robustness to noise, a well-motivated confidence interval for the estimated eigenphase, minuscule memory requirement for the classical control, and speed-up over the standard IPEA.

The potential of RFPE to exhibit these properties has so far been suggested by numerical simulations of the algorithm [16]. Experimental evidence with realistic noise sources is now required to demonstrate the high performance of the approach, potentially vindicating RFPE as a scalable, practical and quadratically faster alternative to other statistical methods [17, 18].

Here we exploit a fully-reconfigurable silicon quantum photonic device, able to generate single photons on-chip and manipulate the states via an arbitrary controlled-unitary operation, to investigate the experimental viability of RFPE. The high level of precision and reconfigurability offered by the photonic chip allow us to compare and contrast the performance of both RFPE and IPEA under the action of different controllable experimental noises. Specifically, we find experimental evidence that RFPE is robust against realistic sources of errors, making it very appealing for near-term useful applications.

Phase estimation— The goal of phase estimation is, given a unitary $\hat{U}$ and a quantum state $|\psi\rangle$, to learn an eigenvalue $e^{i\phi}$ of $\hat{U}$ within the support of $|\psi\rangle$. Standard algorithms work by interfering paths in which either $\hat{\mathbb{1}}$ or $\hat{U}^j$ is applied to $|\psi\rangle$, for integer values of j, and then recombining the paths together to allow them to interfere [3]. Iterative phase estimation works by pooling the results of many such experiments and using a classical inference algorithm to estimate an eigenphase of $\hat{U}$ from the resultant interference pattern.

IPEA allows eigenvalues of $\hat{U}$ to be learned quadratically faster than by statistical sampling and requires exponentially fewer measurements. However, it typically requires long evolution times which can reduce its utility in pre–fault tolerant hardware [12–14]. The most commonly used IPEA algorithm works by inferring each of the bits in a binary expansion of the eigenphase ϕ in reverse order [12–14]. The method uses the circuit in Figure 1a, where the measurement on the control qubit gives output 0 or 1 with probabilities $\cos^2(\pi M[\phi-\theta])$ and $\sin^2(\pi M[\phi-\theta])$, respectively. As each bit is learned iteratively, the algorithm applies a fixed policy for updating

* These authors contributed equally to this work.
† nawiebe@microsoft.com
‡ mark.thompson@bristol.ac.uk

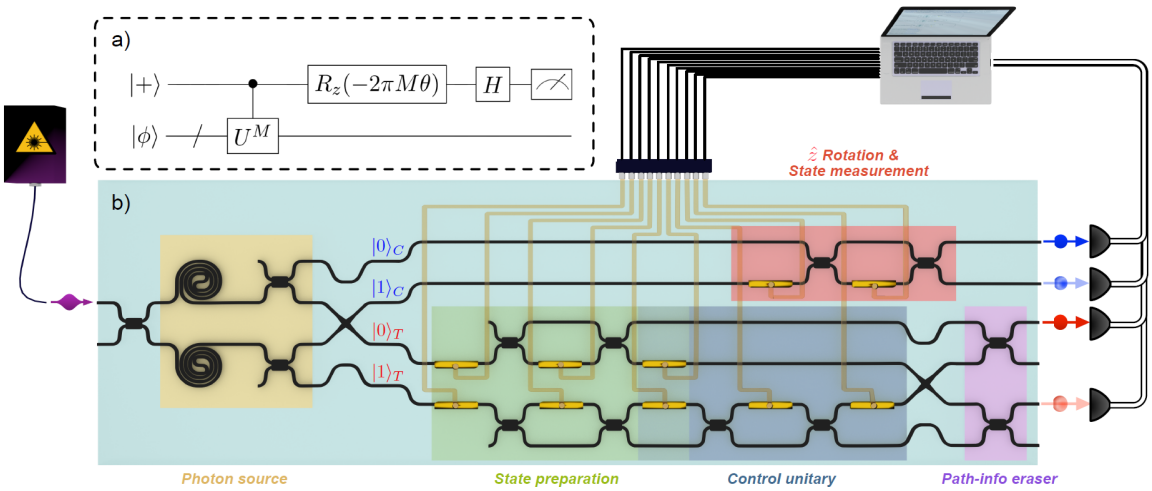

FIG. 1. a). Quantum circuit for standard iterative and Bayesian phase estimation algorithms. b). Experimental setup and the integrated silicon quantum photonic device. The quantum chip can perform any controlled-$U(2)$ operation and any single-qubit state preparation and analysis. Photons are produced and guided in the silicon waveguides (black wires) and reconfigurably controlled by thermo-optical phase shifters. Coherent light was used to generate photons and superconducting nanowire detectors were used for the detection, both coupled to the chip through lensed-single mode fibers. The implementation of the algorithms was achieved by interfacing the quantum device with a classical CPU.

M and θ (see Appendix A).

RFPE is in many ways simpler. This Bayesian approach uses a Gaussian probability distribution $P(\phi)$ (the prior) representing the confidence that the current hypotheses is the correct eigenphase. The result of each new measurement is used to update the mean μ and standard deviation σ according to Bayes' theorem, approximated with rejection sampling (see Appendix B). Specifically, a host of particles are drawn from the prior distribution and then probabilistically discarded based on the likelihood function. The remaining samples model the posterior probability distribution, which becomes the new prior. Since the number of surviving particles decreases exponentially, the posterior distribution is refit to a new Gaussian at each step and fresh particles added, drawn from this new distribution. Rather than learning each of the bits of ϕ individually, the RFPE algorithm gains information about every bit simultaneously.

In more detail, if an outcome $E \in \{0, 1\}$ is obtained from an experiment with parameters M and θ, the likelihood function for the host of particles $\{x_i\}$ is calculated:

$$P(E|x_i; M, \theta)/\kappa = \begin{cases} \cos^2(M\pi(x_i - \theta))/\kappa & E = 0 \\ \sin^2(M\pi(x_i - \theta))/\kappa & E = 1, \end{cases}$$

where κ is a rescaling constant [16]. Particles from the prior are discarded probabilistically based on this function. In order to maximize the information gain at each step, new values of M and θ can be extracted from the prior distribution using various optimization methods. A near-optimal choice is provided by particle guess heuristics [19, 20], giving $\theta \sim P(\phi)$ and $M = \lceil 1.25/\sigma \rceil$. If the likelihood function fails to describe the data (due to experimental noise) then the algorithm estimates the best model for the experimental results within the assumptions, and gives a firm estimate of the uncertainty in the eigenphase. While in standard IPEA any error that occurs in inferring a bit cannot be corrected in subsequent algorithm steps, RFPE does not suffer this issue because it does not try to infer the bits sequentially. While these expectations have been born out in simulation [16], they have not been verified in practice. We provide such verification below using an integrated quantum photonic device. Further details for RFPE and IPEA can be found in Appendix A and Appendix B.

Integrated quantum photonic device — Silicon quantum photonics have emerged as a promising approach for the realisation of quantum hardware, since in principle all the necessary photonic components (sources, circuits, filters and detectors) for quantum information processing can be integrated on a single platform [21]. We developed a quantum photonic device in silicon waveguides for the experimental implementation of RFPE, shown in Fig. 1b. The chip was manufactured on a standard Silicon-on-Insulator platform and is capable of performing arbitrary two-qubit controlled unitary operations. Using spontaneous four-wave mixing, photon pairs were created in two spiral sources pumped with $\simeq$ 20 mW bright light near 1550 nm wavelength [22]. Separating the photons by multi-mode interferometer beam-splitters (BS) and swapping the two modes by a waveguide crosser yield a maximally entangled two-photon post-selected state $(|0\rangle_C |0\rangle_T + |1\rangle_C |1\rangle_T)/\sqrt{2}$, where $|0\rangle$ ($|1\rangle$) indicates the photon state either in its upper or lower spatial mode [23, 24], whereas C (T) subscript refers to the control (target) register. Two additional spatial modes are then added to the target register, obtaining a path encoded qubit for each of the two modes of the target

wave-function. After an initial state preparation $|\phi\rangle_T$, each qubit is manipulated with a separate transformation, depending on which path the photon is travelling on: the identity $\hat{\mathbb{1}}$ for the upper modes (the ones corresponding to a $|0\rangle_C$) and an arbitrary unitary $\hat{V}$ for the bottom modes (the ones corresponding to a $|1\rangle_C$). The path information is erased by two BSs and the state is finally post-selected obtaining the equivalent photon count statistics of an arbitrary control-unitary operation $(|0\rangle_C \otimes |\phi\rangle_T + |1\rangle_C \otimes \hat{V}|\phi\rangle_T)/\sqrt{2}$ [10, 25]. The quantum logic of Fig. 1a is completed by performing a single qubit operation on the control photon. All the processes required for state preparation, manipulation and measurement are achieved through thermo-optical phase shifters and Mach-Zehnder interferometers, as shown in Fig. 1b [24, 26]. Controlling the electric power supplied to the phase shifters, each phase φ can be driven with an average precision of $\simeq 0.01$ rad (see Appendix C and Appendix D). Finally, photons were detected by superconducting nanowire single photon detectors (SNSPD) and coincidence counts obtained by a time interval analyser. The photon statistics were used to measure the projectors $\Pi(0,1)$ on the computational basis of the photonic qubits. The Bayesian update and changes to the controls of the quantum device required by RFPE can be calculated and fed to the quantum system using an interfaced classical computer.

Experimental results — The rapid reconfigurability and the high precision of the silicon photonic device were crucial to the practical implementation and testing of RFPE. As is usual in photonic experiments, where measurements provide probability distributions rather than single-shot data, the value of E was determined by the majority voting method on the projective measurement statistics for both RFPE and IPEA (see Appendix E). For the stochastic representation of the prior, 1000 particles was found to be reliable. The initial prior distribution was set to a Gaussian $\mathcal{N}(\pi, \pi^2)$, which approximates a uniform prior.

In Fig. 2a we report the results from a single RFPE run, which demonstrates that the estimation converges exponentially to the true eigenphase. For this case an error as low as $2.4 \cdot 10^{-4}$ rad is achieved within 50 experiments on the quantum device. This error is in good agreement with the standard deviation of the final posterior Gaussian distribution ($\simeq 4.2 \cdot 10^{-4}$ rad after 50 experimental steps, see the inset of Fig. 2a), thus confirming that the algorithm provides a reliable uncertainty estimate. We remark that this feature is not present in other phase estimation techniques, which do not provide a rigorous estimation of the phase uncertainty that arises in non-fault-tolerant devices [16]. A possible strategy for obtaining such uncertainty estimates from IPEA would be to determine the mean and the standard deviation of the measured eigenphase from repeated experiments. Fig. 2a shows that when adopting this strategy with a reasonable cardinality for the experiments, 10 for the curve reported, RFPE outperforms IPEA.

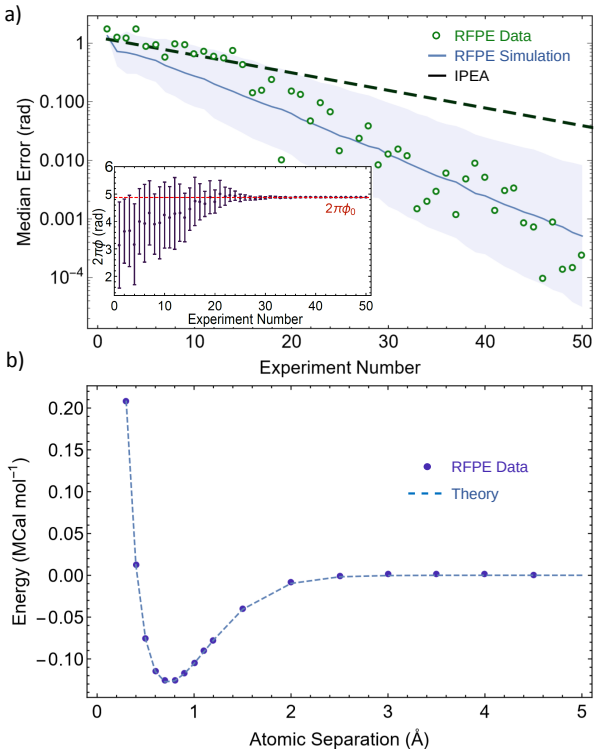

FIG. 2. a) Convergence of RFPE to the true phase value $2\pi\phi_0 = 4.8741$ rad related to the energy of the dissociated H_2 molecule. The initial prior distribution is $\mathcal{N}(\pi, \pi^2)$. Data points show an exponential shrink of the error, compatible with simulations (blue line) of the device performance averaged over 1000 runs of the RFPE algorithm (shaded area: 67.5% credible interval). The dashed black line denotes the convergence of the standard IPEA using the parameters discussed in the main text. Inset: convergence of the phase estimation to ϕ_0 (red line), where errors are evaluated using the s.d. of the prior distribution. b) Bonding energies of the H_2 molecule for various atomic distances using RFPE with 50 steps. Energy estimations are achieved within chemical accuracy. Errors are smaller than the markers and neglected in the plot for more clarity. The dashed line represents theoretical energies.

To extensively test how RFPE performs in key applications, we scanned the energy of molecular H_2 for different atomic separations, using 50 iterations of RFPE for each eigenphase evaluation. The eigenstates of the molecular Hamiltonian are mapped into a the qubit basis using the Jordan-Wigner transformation and the eigenphases are directly related to the binding energy [7, 9–11]. As shown in Fig. 2b, the estimated energies match the theoretical values with high precision. The average error for the dataset is 0.72 kCal/mol, providing a precision higher than chemical accuracy ($\simeq 1$ kCal/mol).

The main advantage of the Bayesian approach over traditional methods comes from its expected reliability on non-fault-tolerant devices. Here we investigate experimentally the robustness of the protocol against two main controllable sources of noises: *gate errors* and *decoher-*

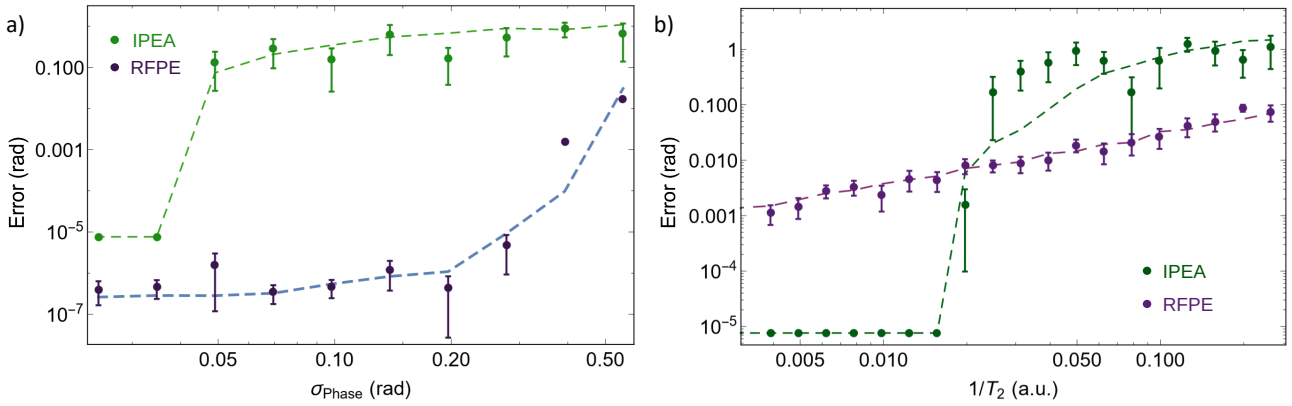

FIG. 3. Effects of different experimental noises on phase estimation strategies. a) Infidelity of quantum operation. Each of the correct phases $\bar{\varphi}_i$ for the phase gates is synthetically replaced with a Gaussian distributed $\varphi_i \sim \mathcal{N}(\bar{\varphi}_i, \sigma_{Phase})$, where σ_{phase} represents a noise in the phases. b) Decoherence. For IPEA data, experiments were repeated 10 times with 16-bit accuracy to evaluate median error and error bars, while the RFPE data were collected from a single run, after 100 measurements, and directly used to evaluate the error and uncertainty within the algorithm. Error bars for the estimated phase represent in both plots either a 67.5% credible region for RFPE, either a 67.5% confidence interval for IPEA. In the cases where error bars are smaller than the markers they have been omitted for clarity. Points are experimental data and dashed lines are simulations averaged over 1000 runs. The simulations take into account the characterized residual phase noise in the device.

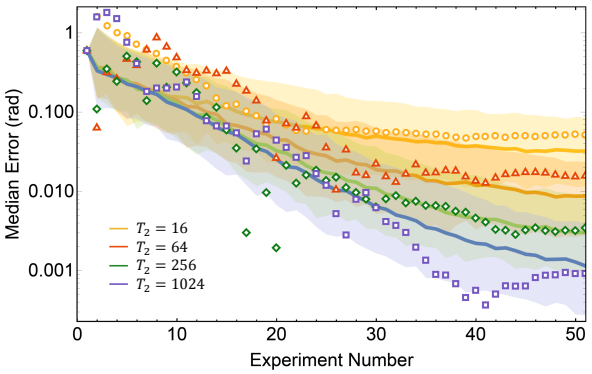

FIG. 4. Experimental data (points) and simulation (lines) of the behaviour of RFPE under different decoherence times T_2. Shaded areas represent a 67.5% credible interval from data simulated averaging over 500 runs of RFPE.

ence.

The infidelity of unitary operations is a well known problem existing in quantum hardware, and is typically given by a noisy control and imperfect manufacture and calibration of the components. On integrated quantum photonic devices it is mainly due to control noise and residual crosstalk of the phase gates, which are expected to occur on the device in Fig. 1 [10]. Our electronic phase-shifters driver allows to precisely control the thermo-optical phase gates as residual thermal cross-talk can be compensated by further calibration (see Appendix D). The high controllability allows us to add a tunable level of noise on the phase gates to study the robustness of RFPE. This is achieved by replacing the correct phases $\bar{\varphi}$ required to implement the unitary transformation with synthetic values φ sampled from a Gaussian distribution $\varphi \sim \mathcal{N}(\bar{\varphi}, \sigma_{\text{phase}})$. The parameter σ_{phase} mimics a Gaussian noise in the phases, which in turn introduces a controllable noise in the fidelity of both the implemented state preparation and the unitary evolution. Fig. 3a shows the convergence of both RFPE and IPEA scanning σ_{phase} up to 0.55 rad, corresponding to average state fidelity 94% and gate fidelity 91% expected in the chip (see Appendix D). We report the performance of RFPE with 100 steps, compared to a 16-bit IPEA, averaged over 10 runs to estimate the error bars, i.e. 160 total experiments per data point. We remark that since a reasonable error bar estimation requires a higher number of experiments for IPEA than RFPE, the relative rates of convergence are not immediately obvious from these figures. Rather, these plots illustrate how the performance of each algorithm is affected by increasing errors, to compare their robustness to noise. For $\sigma_{\text{phase}} \geq 0.05$ rad IPEA dramatically decreases in accuracy and becomes quickly unreliable. This occurs because while the majority voting scheme provides error resilience for small error rates, it can diverge rapidly once the error rate crosses a threshold (see Appendix F).

On the other hand, in this regime the performance of RFPE is initially only slightly affected, maintaining a very high level of precision even when IPEA fails. This is expected because RFPE does not make hard decisions about bits as each experiment yields information about all bits at once. Thus errors are unlikely to be critical. In order for RFPE to be substantially affected we require σ_{phase} higher than 0.3 rad, a value much higher than the actual experimental noise in our device.

Decoherence is an important limitation in many quantum computing experiments but it plays a minor role in integrated quantum photonics. It must then be intro-

duced artificially in our experiment. In order to simulate it, coincidence counts provided by the SNSPDs for the $\Pi(0)$ and $\Pi(1)$ projectors were progressively flattened out by classical post-processing and combined with Poissonian noise in the measurements. In agreement with the depolarizing noise model we mimicked the presence of a normalised decoherence time T_2 (see Appendix G) such that:

$$P'(D|\phi) = e^{-M/T_2} P(D|\phi) + \frac{1 - e^{-M/T_2}}{2},$$

where $P(D|\phi)$ is the data obtained from the photon coincidence counts for the outcome $D \in \{0, 1\}$. We introduce this noise model while processing online the output data during the iterative process, thus affecting the choice of the experiments. In this way it is possible to simulate the behaviour of RFPE and IPEA in systems that are prone to this model of decoherence.

We studied the action of the depolarizing noise up to $T_2 = 4$ for both IPEA and RFPE, as shown in Fig. 3b. The performance of IPEA has a substantial and sharp deterioration at T_2 close to 32, whereas the median error of 100-step RFPE decreases only polynomially with $1/T_2$, maintaining an error $\mathcal{O}(10^{-2})$ even in the regime where conventional IPEA fails to provide any reliable estimate of the phase. In the presence of characterized depolarizing noise an optimized value for M is given by $\min(\lceil 1.25/\sigma \rceil, T_2)$ [16], which, however, implies that, when decoherence is significant, the performance of RFPE degrades significantly. This behaviour is exhibited by the experimental data in Fig. 4, where the convergence of RFPE is reported under the action of various T_2. We observe that RFPE ceases to learn exponentially quickly when $1/\sigma \sim T_2$, after which the algorithm continues to learn at a polynomial rate, unlike IPEA [16].

Conclusion—Our work shows how the precision and controllability developed in quantum technologies, here in particular integrated photonics, allows to go beyond the basic proof-of-principle demonstrations of quantum algorithms and to enter a regime where they can be extensively tested and compared. We experimentally verified the Bayesian phase estimation algorithm on a fully programmable silicon quantum photonic device and demonstrated its superior performance in presence of noise. Although in this work the experiment is performed using a small-scale unitary and a photonic device, more complex future implementation can be efficiently performed on any scalable quantum architecture. The Bayesian approach remarkably lowers the requirements for the implementation of quantum phase estimation on pre-fault-tolerant devices, showing a new promising route for practical and useful quantum information processing in the near future.

Acknowledgements—We thank A. Laing, J.W. Silverstone, D. Bonneau, C. E. Granade and X. Zhou for insightful discussion. We thank K. Ohira, N. Suzuki, H. Yoshida, N. Iizuka and M. Ezaki for the device fabrication. We acknowledge the support from the Engineering and Physical Sciences Research Council (EPSRC), European Research Council (ERC), Photonic Integrated Compound Quantum Encoding (PICQUE), FP7 Action: Beyond the Barriers of Optical Integration (BBOI), Quantum Simulation on a Photonic Chip (QuChip), US Army Research Office (ARO), and Centre for Nanoscience and Quantum Information (NSQI). D.P.T thanks the Royal Society for a University Research Fellowship. J.L.O'B. acknowledges a Royal Society Wolfson Merit Award and a Royal Academy of Engineering Chair in Emerging Technologies. M.G.T. acknowledges the support from an EPSRC Early Career Fellowship and the Toshiba Research Fellowship scheme.

[1] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," Proc. 35nd Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press (1994).

[2] S. Lloyd, "Universal quantum simulators," Science **273**, 1073–1078 (1996).

[3] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2001).

[4] B. P. Lanyon *et al.*, "Experimental demonstration of a compiled version of shor's algorithm with quantum entanglement," Phys. Rev. Lett. **99**, 250505 (2007).

[5] E. Martin-Lopez *et al.*, "Experimental realization of shor's quantum factoring algorithm using qubit recycling," Nature Phot. **6**, 773 (2012).

[6] T. Monz *et al.*, "Realization of a scalable shor algorithm," Science **351**, 1068–1070 (2016).

[7] B. P. Lanyon *et al.*, "Towards quantum chemistry on a quantum computer," Nature Chemistry **2**, 106–111 (2010).

[8] J. Du, N. Xu, X. Peng, P. Wang, S. Wu, and D. Lu, "Nmr implementation of a molecular hydrogen quantum simulation with adiabatic state preparation," Phys. Rev. Lett. **104**, 030502 (2010).

[9] P. J. J. O'Malley, R. Babbush, *et al.*, "Scalable quantum simulation of molecular energies," Phys. Rev. X **6**, 031007 (2016).

[10] R. Santagati, J. Wang, *et al.*, "Quantum simulation of hamiltonian spectra on a silicon chip," arXiv:1611.03511v3 .

[11] A. Aspuru-Guzik, A. Dutoi, P. J. Love, and M. Head-Gordon, "Simulated quantum computation of molecular energies," Science **309**, 1704–1707 (2005).

[12] A. Y. Kitaev, "Quantum measurements and the abelian stabilizer problem," Electronic Colloquium on Computational Complexity **3** (1996).

[13] R. B. Griffiths and C.-S. Niu, "Semiclassical fourier transform for quantum computation," Phys. Rev. Lett. **76**, 3228–3231 (1996).

[14] M. Dobšíček, G. Johansson, V. Shumeiko, and

G. Wendin, "Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark," Phys. Rev. A **76**, 030306 (2007).
[15] M. A. Rubin and S. Kaushik, "Loss-induced limits to phase measurement precision with maximally entangled states," Phys. Rev. A **75**, 053805 (2007).
[16] N. Wiebe and C. Granade, "Efficient bayesian phase estimation," Phys. Rev. Lett. **117**, 010503 (2016).
[17] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," Nature Commun. **5**, 4213 (2014).
[18] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," New Journal of Physics **18**, 023023 (2016).
[19] C. Ferrie, C. Granade, and D. G. Cory, "How to best sample a periodic probability distribution, or on the accuracy of hamiltonian finding strategies," Quantum Information Processing **12**, 611–623 (2013).
[20] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, "Hamiltonian learning and certification using quantum resources," Phys. Rev. Lett. **112**, 190501 (2014).
[21] J. Silverstone, D. Bonneau, J. L. O'Brien, and M. G. Thompson, "Silicon quantum photonics," IEEE **22**, 6700113 (2016).
[22] M. D. Reid and D. F. Walls, "Quantum theory of non-degenerate four-wave mixing," Phys. Rev. A **34**, 4929 (1986).
[23] J. Wang *et al.*, "Chip-to-chip quantum photonic interconnect by path-polarization interconversion," Optica **3**, 407–413 (2016).
[24] J. W. Silverstone, R. Santagati, D. Bonneau, M. J. Strain, M. Sorel, J. L. O'Brien, and M. G. Thompson, "Qubit entanglement between ring-resonator photon-pair sources on a silicon chip," Nature Commun. **6**, 7948 (2015).
[25] X. Zhou, P. Kalasuwan, T. C. Ralph, and J. L. O'Brien, "Calculating unknown eigenvalues with a quantum algorithm," Nature Photon. **7**, 223–228 (2013).
[26] G T Reed, G Mashanovich, F Y Gardes, and D J Thomson, "Silicon optical modulators," Nature Photon. **4**, 518–526 (2010).
[27] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O'Brien, "Silica-on-silicon waveguide quantum circuits," Science **320**, 646–649 (2008).
[28] M. Veldhorst *et al.*, "A two-qubit logic gate in silicon," Nature **526**, 410 (2015).

Appendix A: Iterative Phase Estimation Algorithm (IPEA)

In the standard implementation of IPEA less significant digits of the unknown phase are evaluated first using the circuit in Fig. 1A. The information obtained by the measurement is transferred to the quantum circuit and allows to refine the measurement of more significant digits via an R_z rotation on an ancillary qubit. This procedure increases the probability of obtaining the correct phase value. If the phase has a binary expansion with a finite number of digits, say n bits, the algorithm is deterministic and requires n iterations of the experiment.

Each digit of the eigenphase is iteratively inferred, starting from the least significant one. The exponent M is chosen so that, if we want to obtain ϕ with n bits precision, at the j-th iteration we have $M_j = 2^{n-j}$. Supposing that ϕ has the finite binary expansion $\phi = 0.\phi_1 \dots \phi_n$, the first iteration step is performed using $\theta_1 = 0$ and $M_1 = 2^{n-1}$, which from Eq. B2 and Eq. B3 results in $P(0) = \cos^2(\pi\ 0.\phi_n)$. The data obtained from the measurement is 0 if $\phi_n = 0$ and 1 if $\phi_n = 1$, inferring the last bit of ϕ deterministically. For the j-th iteration we choose $\theta_j = 0.0\phi_{n-j+2} \dots \phi_{n-1}\phi_n$ so that, with a same reasoning as before, the digit ϕ_{n-j+1} is exactly determined from the projective measurement of the ancilla. Iterating this process n times we can reconstruct the whole binary expansion of ϕ.

In presence of experimental noises, the probability of measuring the correct value of the k-th bit for single-shot measurements, given in Ref.[14], is

$$P_k(\Delta_x, T_2) = \frac{\left(1 + e^{-\Delta_x^2 - a2^k T_2}\right)}{2}, \tag{A1}$$

where the errors considered here are a dephasing with decoherence time T_2 and $R_x(\delta)$ errors on the control qubit for the state preparation and measurement projection, with the angle δ_x Gaussianly distributed with variance Δ_x. a is a system depended parameter [14]. The susceptibility of the algorithm to experimental noise is evident from Eq.A1. The success probability can be increased repeating the measurements for each qubit and adopting a majority voting scheme, as we do in our implementation, but still maintains an exponential scaling with the noise level. Moreover, if an error occurs, IPEA is not able to correct it and has limited capability to recover from it. Therefore, as noise become significant, errors cumulate and make the algorithm fail.

Appendix B: Rejection Filtering Phase Estimation (RFPE)

We describe the Bayesian phase estimation algorithm, RFPE, introduced in [16] and employed in our experiment. At a fundamental level, RFPE involves the exact same experimental setups as traditional iterative phase estimation. The main differences arise in the way RFPE chooses experiments and also how the experimental data is processed.

The fundamental object in RFPE, as well as all other Bayesian approaches, is the prior distribution $P(\phi)$. It describes the algorithms confidence that a particular eigenphase $\phi \in [0, 2\pi)$ corresponds to the eigenvector provided to the algorithm. As experiments are performed, the prior distribution is then updated based on the likelihood of the observed data.

Bayes' theorem provides the correct way to update such prior beliefs when provided with evidence. The theorem states that if an event E is observed for an experiment parametrized by M and θ then the updated prior distribution should be

$$P(\phi|E; M, \theta) = \frac{P(E|\phi; M, \theta)P(\phi)}{\int P(E|\phi; M, \theta)P(\phi)d\phi}, \tag{B1}$$

where $P(E|\phi; M, \theta)$ is known as the likelihood function and the updated prior distribution $P(\phi|E; M, \theta)$ is known as the posterior distribution.

The central challenge of Bayesian phase estimation is that even if the prior distribution is discretized then the complexity of storing the prior distribution to sufficient accuracy can be immense. For example, if error on the order of 10^{-9} is needed, then several gigabytes of memory will be needed to store the distribution. This makes it impractical for use in Shor's algorithm. Also, if the data arises from more than one eigenvalue then it is necessary to store multiple hypotheses. This exponentially increases the number of hypotheses that need to be tracked, rendering such Bayesian inference impractical.

RFPE avoids this problem by invoking a form of assumed density filtering that is based on rejection sampling. This approach allows inference to be performed quickly in a memory limited environment without the need to compute costly functions. The algorithm works by positing a model for the prior distribution, which for convenience is usually taken to be a Gaussian. The goal of each update is then to find the best model for the posterior distribution within the family of allowed models. The procedure involves drawing a sample from the prior distribution and then computing

the likelihood that the data would emerge if that eigenphase were correct. The sample is then randomly accepted with probability equal to this likelihood. The mean and covariance matrix of these accepted samples can easily be shown to match the posterior mean and covariance. Therefore these statistics provide the best model for the posterior distribution within any family of distributions that are parametrized by their first two moments. This approach is therefore well suited for experiment because it naturally outputs an estimation for its uncertainty in the form of the covariance matrix, or variance, of the distribution.

The likelihoods used in this algorithm correspond to the measurement probabilities predicted by quantum mechanics for the circuit in Fig. 1A of the main text. Specifically, if an experiment with parameters (M, θ) is performed, the measurement on the ancillary qubit will have outcome $E \in \{0, 1\}$ with probability:

$$P(0|\phi; M, \theta) = \cos^2(\pi M[\phi - \theta]), \tag{B2}$$

$$P(1|\phi; M, \theta) = \sin^2(\pi M[\phi - \theta]). \tag{B3}$$

An important remaining issue is that of choosing the parameters M and θ that, given the prior distribution, optimize the information gain from the experiment. In our experiment we have used the Particle Guess Heuristic (PGH) [19, 20], which provides near-optimal values and does not require any precomputation. In this approach the experiment is given by $M = \lceil 1.25/\sigma \rceil$, $\theta \sim \mathcal{N}(\mu, \sigma)$. The fact that these parameters depend intimately on the current estimate of the uncertainty means that RFPE is implicitly an adaptive protocol. Furthermore, the fact that RFPE uses prior information means that the asymptotic optimality result of [19], which is based on the Cramer-Rao bound, does not apply in this setting. It was further shown that this heuristic approximately saturates the Bayesian Cramer-Rao bound [16] under the assumption of a Gaussian prior and thereby demonstrating that the PGH is asymptotically near-optimal under these assumptions.

The full RFPE algorithm is described in Algorithm 2. For further details regarding, for example, the number of particles required to achieve small errors and variance reduction strategies, we refer to [16].

Algorithm 1. Bayesian update via Rejection Filtering

Input: experimental data E, mean μ and variance σ of the prior distribution, experiment parameters M and θ, number of particles $N_{\mathrm{P}art}$, scale κ_E.
Output: mean μ' and variance σ' of the posterior distribution.

function UPDATE$_{\mathrm{RF}}(E, \mu, \sigma, M, \theta, N_{\mathrm{P}art}, \kappa_E)$
  $\mu_{\mathrm{acc}}, \mu'_{\mathrm{acc}}, V_{\mathrm{acc}}, V'_{\mathrm{acc}}, N_{\mathrm{acc}} = 0.$ ▷ Initialize accumulators.
  for $i \in 1 \to N_{\mathrm{P}art}$ **do**
    $x \sim \mathcal{N}(\mu, \sigma).$ ▷ Sample a particle according to prior distribution.
    $x = x \bmod 2\pi.$
    $x' = x + \pi \bmod 2\pi.$
    $u \sim \mathrm{Uniform}(0, 1).$ ▷ Accept the particle with probability $P(E|x)$.
    if $P(E|x) \geq \kappa_E u$ **then**
      $\mu_{\mathrm{acc}} = \mu_{\mathrm{acc}} + x.$
      $V_{\mathrm{acc}} = V_{\mathrm{acc}} + x^2.$
      $V'_{\mathrm{acc}} = V_{\mathrm{acc}} + x'^2.$
      $N_{\mathrm{acc}} = N_{\mathrm{acc}} + 1.$
    end if
  end for
  $\mu' = \mu_{\mathrm{acc}}/N_{\mathrm{acc}}.$ ▷ Calculate mean and variance of the posterior distribution using accepted particles.
  $\sigma' = \min\left(\sqrt{(V_{\mathrm{acc}} - \mu^2_{\mathrm{acc}})/(N_{\mathrm{acc}} - 1)}, \sqrt{(V'_{\mathrm{acc}} - \mu^2_{\mathrm{acc}})/(N_{\mathrm{acc}} - 1)}\right).$
  return $(\mu', \sigma').$
end function

Algorithm 2. Rejection Filtering Phase Estimation (RFPE)

Input: Initial prior distribution $\mathcal{N}(\mu_0, \sigma_0)$, total number of experiments N_{Steps}, number of particles N_{Part}, scale κ_E.
Output: the phase estimation μ and its uncertainty σ.

$\mu = \mu_0$, $\sigma = \sigma_0$. ▷ Initialize the mean μ and variance σ of the distribution.
for $i \in 1 \to N_{Steps}$ **do**
  $M = \lceil 1.25/\sigma \rceil$, $\theta \sim \mathcal{N}(\mu, \sigma)$. ▷ Get approximate optimal parameters M and θ for the experiment via PGH.
  Get data E from the experiment using parameters M and θ.
  $(\mu, \sigma) = \text{UPDATE}_{RF}(E, \mu, \sigma, M, \theta, N_{Part}, \kappa_E)$. ▷ Perform a Bayesian update via Rejection Filtering.
end for
return (μ, σ).

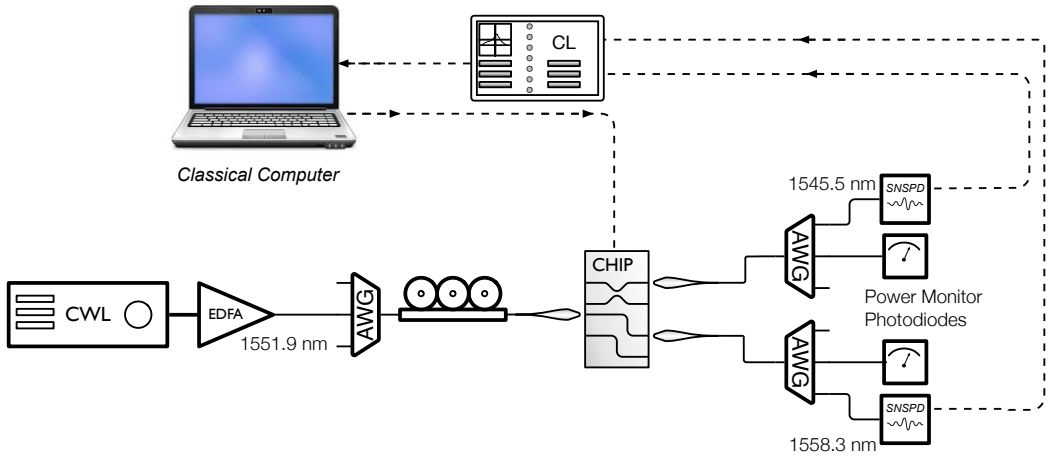

FIG. 5. Schematic representation of the experimental set-up. The light from a CW laser light source is amplified through an EDFA. The pump is filtered through an AWG and the polarisation adjusted through a manual polarisation controller, then focused inside the chip and collected by lensed-fibres. The emerging light is then filtered through AWG, where the single photons channel and the the pump are separated. While the CW light is detected by standard photo-diodes for monitoring the coupling, the single photons are detected through super-conductive nano-wire single photon detectors SNSPDs, whose pulses are analysed by a time interval analyser connected to a classical computer. The classical computer is also used to control the on-chip thermo-optic phase shifters.

Appendix C: Experimental setup

1. The silicon device and experimental details

The device is a silicon-on-insulator chip fabricated on a silicon-on-insulator material system using 248 nm deep-UV photolithography and dry etching. Light is horizontally coupled into and out of the device via spot-size converters with 300 μm long inverse taper with a 200 nm wide tip beneath a 4×4 μm^2 polyimide waveguide. Approximately 8 dB coupling loss per facet was observed between the chip and a single-mode lensed fiber with a 3 μm mode-diameter. Inside the chip, the photons propagate in integrated single-mode waveguides, designed with a width of 450 nm and thickness of 220 nm and covered with a 1 μm silicon dioxide upper cladding. Multi-mode interferometers (MMIs) with a footprint of 2.8 $\mu m \times 28$ μm were used to realize integrated beam-splitters with near 50% reflectivity.

To stabilise the temperature of the silicon device, the chip was mounted and wire-bonded on a PCB that was glued by thermal hypoxy on a Peltier device connected to a heat-sink. The temperature was controlled by a PID algorithm.

The input/output optical fibers were automatically recoupled before each scan of the algorithms to maximize the coupling efficiency between waveguides and fibers.

The device integrates key functionalities for photonic quantum information processing, such as sources to generate an entangled photon-pair and arbitrary single- and two-qubit gates. The photons are generated in two spiral waveguide sources via spontaneous four-wave mixing (SFWM), achieved exploiting the intrinsic $\chi_{(3)}$ non-linearity in silicon [23,

24]. The two spirally-wrapped waveguide sources were designed with a 1.2 cm length to enhance the photon-pair generation. In the device the qubits are realized encoding the information in the path of the photons [27]. The sources were pumped using a continuous-wave (CW) bright laser at 1551.9 nm, amplified by an erbium doped fiber amplifier (EDFA), with an emerging power of $\sim 10\ mW$. Using a fiber-polarization controller before the silicon device we ensured that the transverse-electric (TE) polarised CW light was injected into the device, after the background was removed through the use of array waveguide gratings wavelength division multiplexers (WDM).

Equivalently to the scheme presented in [23], the two coherently pumped sources generate a pair of path-entangled photons, i.e. *signal* and *idler* photons, via the SFWM process. The wavelengths of the signal and idler photons are selected at 1545.5 nm and 1558.3 nm, respectively. Considering the use in this experiment of non resonant SFWM photon pair sources and of a CW laser, the generation probability of multi-photon-pair events is negligible. The output photons, after evolving through the integrated device, were coupled into single-mode fibers, separated apart from the pump light by off-chip AWG with a 0.9 nm bandwidth and $> 90\ dB$ extinction. Photons were finally detected using superconducting PhotonSpot$^{\text{TM}}$ nano-wires single-photon detectors (SNSPDs), with system detection efficiency higher than 85%, sub-100Hz dark counts, approximately 70 ps FWHM timing jitter and recovery times of approximately 50 ns. The coincidence counts were recorded by the use of a time interval analyser (Picoharp 300 by PicoQuant$^{\text{TM}}$) with an integration time of 10 s per data point and a coincidence window of approximately 400 ps. A maximal photon-pair rate of approximately (200 Hz) was observed.

The active and reconfigurable quantum operations inside the silicon device rely on thermo-optical phase shifters, which are formed by metal resistive heaters on top of the silicon waveguides isolated by a dioxide layer in between. The heat locally dissipated in the waveguide induces the refractive index change responsible of the phase-shift. The heaters can be independently driven and controlled by home-designed multichannel electronics interfaced with a classical computer. This allows to control the phase with a precision of approximately 0.01 rad, as reported in more details later on.

2. State evolution and scheme for arbitrary controlled-unitaries

The logic circuit for the iterative phase estimation algorithms, both RFPE and IPEA, requires a controlled-unitary between the two qubits, as shown in Fig. 1a of the main text. In this section we describe in more details the entanglement-based scheme we used to obtain arbitrary controlled operation in the quantum photonic device.
Coherently pumping the two spiral SFWM sources, a pair of photons is generated in a superposition between them. Using a Fock representation, the photon number state generated is written as $(|0200\rangle + |0020\rangle)/\sqrt{2}$. Two MMIs beam-splitters following the sources probabilistically split the photons. Post-selecting the cases where the signal photon is collected from the two upper modes and the idler from the four bottom modes, analogously to the scheme shown in [10, 23, 25], the state becomes:

$$\frac{|1010\rangle + |0101\rangle}{\sqrt{2}}. \tag{C1}$$

To encode the qubit in the photons' path, for each photon we use the convention $|0\rangle \leftrightarrow |10\rangle$ and $|1\rangle \leftrightarrow |01\rangle$ to convert the number state to the logic state. Using this convention, the state obtained from the sources can be rewritten as the maximally-entangled Bell state $(|0\rangle_C|0\rangle_P + |1\rangle_C|1\rangle_P)/\sqrt{2}$.
Inserting an additional degree of freedom to the path of the idler photon, we obtain the entangled state $(|0\rangle_C|0\rangle_T|0\rangle_P + |1\rangle_C|1\rangle_T|1\rangle_P)/\sqrt{2}$. The two components associated to the second qubit go through one R_z and one R_y gate to prepare the qubit in the same state $|\psi\rangle_T$ for both pairs of paths. These two components then go through two different operations: $\hat{I}$ for the upper pair of modes in the target register and $\hat{U}$ for the bottom ones, obtaining an entangled state of the form:

$$\frac{|0\rangle_C|\psi\rangle_T|0\rangle_P + |1\rangle_C(\hat{U}|\psi\rangle_T)|1\rangle_P}{\sqrt{2}}. \tag{C2}$$

To obtain a superposition of these two different operations, we erase the path information between the two components of the target state. This is obtained employing two waveguide crossings and combining the modes in two MMI beam-splitters, giving

$$\frac{(|0\rangle_C|\psi\rangle_T + |1\rangle_C\hat{U}|\psi\rangle_T)|0\rangle_P + (|0\rangle_C|\psi\rangle_T - |1\rangle_C\hat{U}|\psi\rangle_T)|1\rangle_P}{2}. \tag{C3}$$

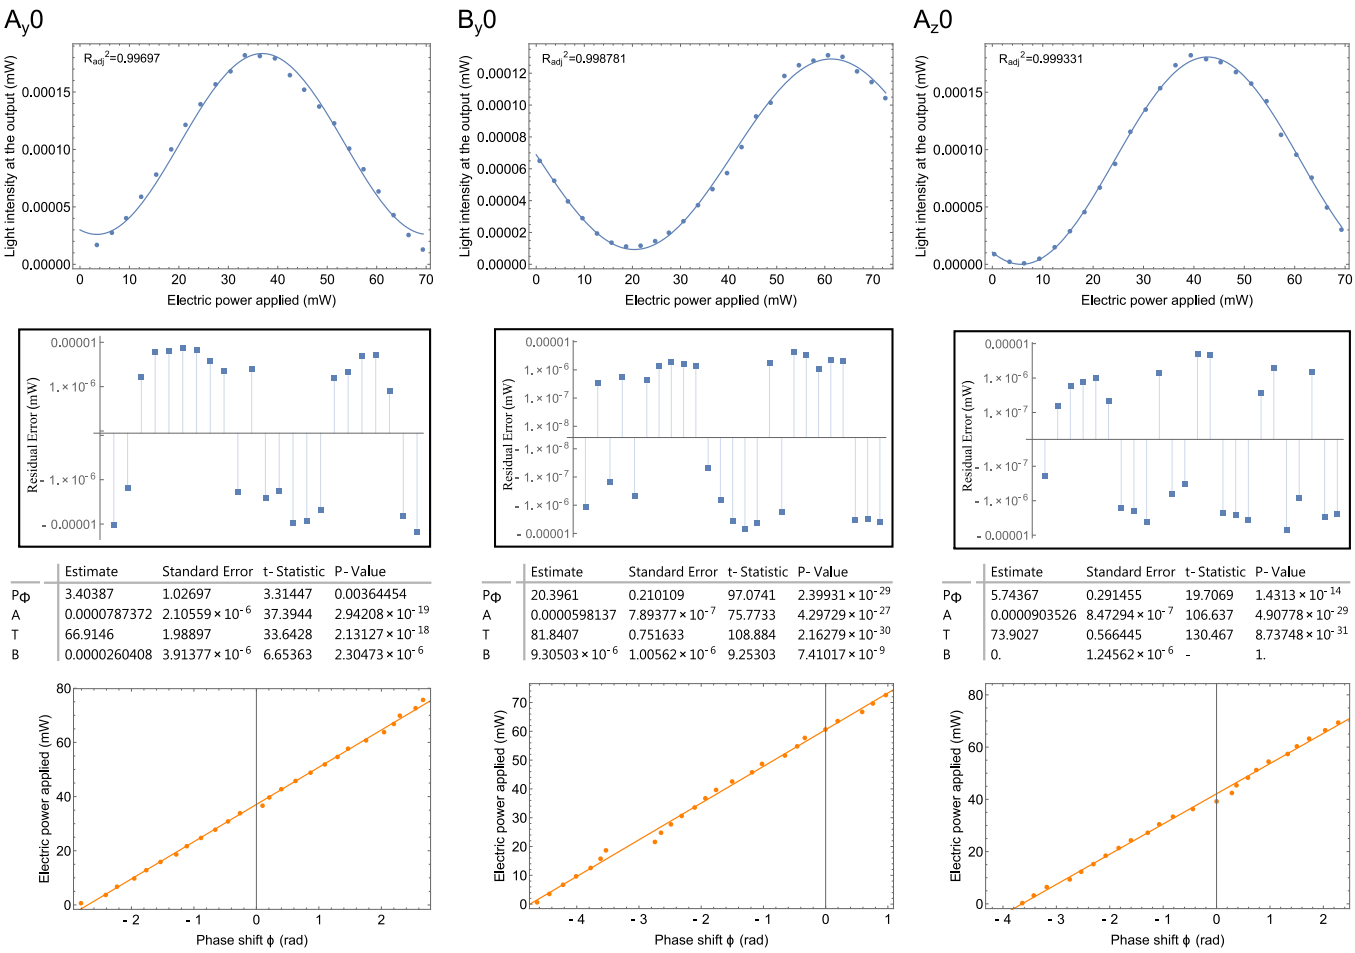

FIG. 6. Characterization for the implemented phase on three exemplary phase shiifters on the device, chosen as a worst case (A_y0), an average case (B_y0) and a best case (A_z0). For each phase shifter, we report in column, from top to bottom: i) the nonlinear fit of raw experimental data, collected applying electrical power to the phase shifter, while measuring the optical power outputted at the corresponding optical port (vertical axis). The fitting is performed using the nonlinear Eq. D1. The value of R^2 is reported in the plot for each separate case. ii) A plot of the residuals corresponding to the nonlinear fits mentioned before. iii) A table of the parameters obtained in the same nonlinear fit is displayed. For each free parameter in the nonlinear model, we report an estimate of both the fitted value, as well as its statistical uncertainty (Standard Error), and a test of its significance (t-Statistic and P-value). iv) For clarity, we report also the linear dependence of the phase shift implemented by each heater against the corresponding electrical power applied.

Finally, projecting the third qubit into $|0\rangle_P$, the final state can be represented as

$$\frac{|0\rangle_C|\psi\rangle_T + |1\rangle_C\hat{U}|\psi\rangle_T}{\sqrt{2}} \tag{C4}$$

which is equivalent to apply the desired arbitrary control-unitary operation.

We remark that this is a post-selected scheme, which makes it not scalable. However, it enables a wide range of experiment previously inaccessible to integrated quantum photonics.

Appendix D: Estimating and simulating phase errors in the integrated photonic device

The quantum photonic experimental setup used in this experiment adopts thermo-optical phase shifters to implement both the quantum state preparation and unitary evolution of the qubits. The phase shifters are integrated on-chip as metallic thin-film tracks (depicted with golden rods in Fig. 1B), coating the top of the silicon waveguides. When supplied with DC current, the phase shifters act as non-Ohmic resistive heaters, thus changing locally the

refractive index of the waveguide [26]. Each of the heaters is driven and controlled independently using multi-channel current drivers, in order to reduce electrical cross-talk among the heaters. Cross-talk is expected because of a shared ground connection in the quantum photonic circuit design. Each driver was hosted in a board featuring 12 bits DAC, controlled by the experiment software via RS-232 interface.

Thermal cross-talk among the phase shifters may also occur in a thermo-optically controlled device, because of heat flowing across the dioxide layer and the electrical connections to waveguides other than the targeted one. Thermally induced, unwanted changes in the refractive index of other optical paths alter the implemented phase shifts in a non-controllable way, introducing systematic errors in the setup. In order to compensate for this effect, different calibrations for different chip configurations were preliminary run. For each configuration, all the heaters are constantly driven at a certain current, set to implement an appropriate state preparation and/or gate operation, except the heater to be calibrated. This heater is supplied with a range of currents (and thus electrical power, P_{el}), as oscillations in optical power P_{op} at the output of the corresponding optical path are recorded. The collected data can be fitted according to the non-linear function

$$P_{op} = B + A\cos[\frac{2\pi}{T}(P_{el} - P_{\Phi})] \tag{D1}$$

where B is a background, A is the maximum amplitude of the signal, T its period, and P_{Φ} is the offset power value for the heater in such a configuration, all obtained from the model fit. Detailed calibrations are reported in Suppl.Fig.6 for three examplary phase shifters, where we also report the relevant statistical parameters of each fit. The fits were performed using the built-in functions from *Mathematica* ®, that estimate the parameters using a (non-linear) least-squares approach, where the optimization is performed using a numerical global optimization method ("RandomSearch"). All fits show R^2 values close to one, thus suggesting that the model can adequately reproduce the data observed from the measurements. Also, we report high t-statistics and low p-values for most of the parameters, testing the significance of all the parameters adopted in the model, with the exception of the aforementioned background B parameter in those cases where the fringe visibility is particularly high (i.e. no significant background is present). In conclusion, the data analysis gives evidence of the suitability of the model in Eq. D1 to describe the physics of our device.

After the calibration, the targeted phase $\bar{\varphi}$ for a heater is obtained driving the heater with P_{el} such that:

$$\bar{\varphi} = \frac{2\pi}{T}\,(P_{el} - P_{\Phi}) \tag{D2}$$

and therefore, propagation of stochastic errors in P_{el}, T, P_{Φ} affects $\bar{\varphi}$. From experimental non linear fits, relative statistical uncertainties are $\sigma_{P_{\Phi}}/P_{\max} \simeq 0.2\%$ and $\sigma_T/T \simeq 1.1\%$. Therefore, inaccuracies in the current supplied by the driver to each heater ($\pm 0.005\ mA$), affecting the actual value of P_{el}, can be neglected as they are less than 0.04% for all the heaters, in the standard configuration used for phase estimation experiments in this paper.

Confining the propagation of errors to only T, P_{Φ}, and averaging over the full interval of P_{el} adopted in the experiment (appr. $5-80\ mW$, slightly different for each heater due to calibration), an average precision $\sigma_{exp} \simeq 0.01\ rad$ can be estimated for the phases, as experimentally implemented by the heaters. In conclusion, systematic errors in the setup can be drastically reduced via an accurate calibration procedure. Nevertheless, in our experiment the effective phase φ implemented by each of the phase shifters is affected intrinsically by a stochastic uncertainty σ_{exp}. This can be estimated from the fits of experimental data. A natural way to investigate the role of imperfect calibration in our device is then to synthetically replace each of the $\bar{\varphi}$, implementing the correct state and unitary preparation, with a phase value sampled from a Gaussian distribution characterized by a variance $\sigma_{phase} > \sigma_{exp}$. An increase in the stochastic noise introduced (σ_{phase}) corresponds to a decrease in the fidelities for state preparation and gate implementation. A simulation of this dependency is reported in Fig. 7

Appendix E: Photonic implementation of RFPE

In the original proposals both IPEA and RFPE algorithms require a single-shot measurement from the experiment [12, 14, 16]. However, in quantum photonics and other quantum platforms, e.g. diamond NV^- centres, measurements usually return a number N of samples distributed according to the output distribution rather than single-shot data, where N depends linearly on the integration time for the data collections. In our experiment photon counts were collected along $10s$ integration time, leading to typical total counts of ~ 2000 for each measure, well beyond the single-shot case. These types of measurements are clearly more informative than the single-shot ones. Nevertheless, phase estimation algorithm (and particularly adaptive implementations) are designed to extract information from single measurements. A challenge in optical implementations is then to convert bulk measures into single-shot data, effectively preserving their enhanced information.

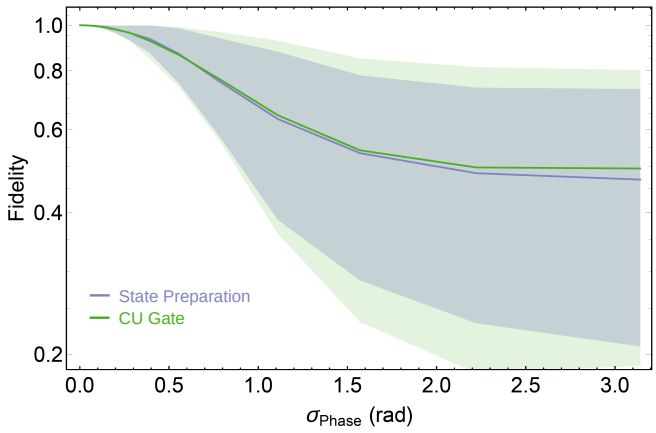

FIG. 7. Simulation of the quantum photonic device performances in terms of the average state preparation fidelity (violet) and CU-gate fidelity (green) for different levels of Gaussian noise in the phase shifters, as described by σ_{phase} spanning the full interval of values tested in the experiments reported in the main paper. The behaviour for the simulated fidelity of the two operations is very similar, with small discrepancies.

In photonic experiments the standard way to extract single-shot data from bulk measurements of a qubit is to classically collapse the corresponding binomial distribution adopting a majority voting scheme, thus deciding each single-shot datum according to the most frequent outcome. That is, if the number of photons measured in the state $|0\rangle$ is higher that those in $|1\rangle$ then the output data is $D = 0$, otherwise it is $D = 1$. This method makes use of some of the information contained in the bulk measurement, enhancing the stability and robustness to experimental noise of the algorithm subtending the measurements. Historically all photonic implementations of IPEA, including ours, have been performed using this scheme [7, 9, 10].

While the advantages of using a majority voting scheme are clearly maintained for phase estimation protocol, the validity of the approach for RFPE is more subtle. In fact, a majority voting scheme does not produce single-shot data $D \in \{0, 1\}$ distributed according to the distribution $\Pr(D|\phi; M, \theta) = \cos^2(\pi M[\phi - \theta] + D\pi/2)$, thus making likely suboptimal the RFPE algorithm, which adopts the likelihood function $\Pr(D|\phi; M, \theta)$. For this reason, the adoption of a majority voting scheme can be potentially problematic. When estimating phases from a photonic device, it slows down the inference process because of inefficiencies in the information extracted. When instead the task is to benchmark photonic implementations against experimental platforms based upon single-shot measurements, a majority voting introduces biases. This is because each run always observes the most likely measurement outcome and feeds it into the learning process, instead of reproducing the measurement statistics expected from random single instances. On one side, this reduces the possibility to retrieve contradictory results from the measurements, speeding up the inference. At the same time, however, it handicaps adaptive algorithms like RFPE, that are optimized for the case when at each step a single measure is sampled from the likelihood function.

It is therefore important to compare majority voting with other strategies that could be used to implement RFPE in a photonic environment, whose measurements reproduce better random samples from the true likelihood function, as ideally expected from the algorithm. Results are reported in Fig. 8.
As an alternative approach, we considered to naïvely sample N times the value for each single datum from the measured statistics of the photons. This effectively resembles the true likelihood function $\Pr(D|\phi; M, \theta) = \Pr(D|\phi_{\text{true}}; M, \theta)$. For a high number of photons this is equivalent to sampling the single-shot data from the true distribution of the outputs (in the absence of systematic experimental errors), which is equivalent to the likelihood $P(D|\phi_{\text{true}}; M, \theta)$. The case where one only photon is sampled randomly from this distribution then coincides with the assumed likelihood function given in the main body for RFPE. The case where a very large number of photons is sampled corresponds to inferring on the actual data observed, which we process within this framework by using RFPE iteratively on N samples drawn from the empirical likelihood. Cases that process many sampled photons are not expected to inherit the asymptotic optimality given by the particle guess heuristic adopted in RFPE [16, 19].

We compare the performance of these approaches in Fig. 8. The first thing to note is that majority voting outperforms the $N = 1$ version of RFPE, requiring polynomially fewer measurements to achieve the same fixed accuracy. This is expected: the inference process in RFPE benefits here from data whose uncertainties are greatly reduced compared to the $N = 1$ case.

On the other hand, we observe that small numbers of sampled measurements $N = 2, 3, 4, 5$ all outperform majority voting. This indicates how already a few random samples give access to a richer information than provided by majority

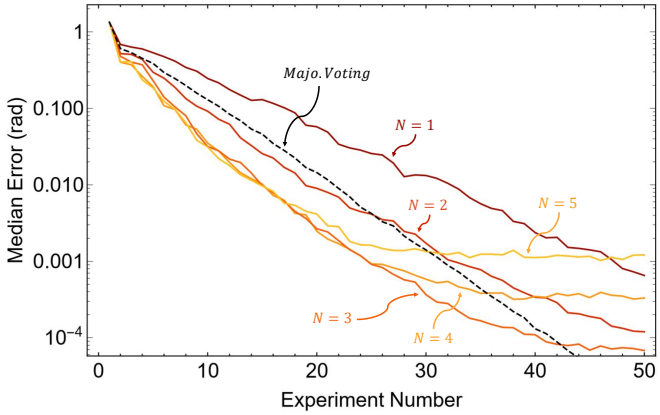

FIG. 8. Performance of RFPE using two different strategies: sampling N values from the output distribution measured or using majority voting. The case $N = 1$ corresponds to single-shot measurements. Simulated data are averaged over 200 runs.

voting at the beginning of the learning process. However, examining later times it is evident how for these same cases RFPE fails to learn in the median for these records. This diminished robustness is results from taking aggressive experiments in RFPE that lead to posteriors that violate the Gaussian assumptions. It can be corrected by taking shorter experiments or using the variance reduction strategies in [16]; however, both constitute adapting RFPE to better process data that naturally arise from photonics experiments

To conclude, these results suggest that though the majority voting approaches commonly used in photonics experiments lead to data records that are atypical of those expected from many quantum computing experiments ($N = 1$ case), majority voting provides qualitatively similar results while partially exploiting the enhanced information extraction of photonic measurements. This, together with the facility to make the comparison with IPEA via the same measurement approach, justifies the use of majority voting for the scope of this work. Finally, considering how all the strategies compared in this paragraph do not fully exploit the entire predictive power of the observed data records and the power of the Bayesian approach implies that there is still space to improve RFPE performances for quantum photonics platform. If we wish to fully assess the power of photonic quantum devices, we need to modify algorithms like RFPE to use likelihood functions that are appropriate for the bulk measurements seen in photonics experiments and transition away from using majority voting as a surrogate for performing proper statistical inference on the data.

Appendix F: Analysis of Breakdown of Majority Voting in IPEA

To provide an understanding of why IPEA breaks down so dramatically in the presence of substantial noise, it is instructive to examine a simple model. Here we consider the model of depolarizing noise, which is similar to that of T_2 we considered but for convenience we assume that the strength of the noise is independent of the evolution time. In particular, let P be the probability of measuring 1 in an experiment and let us assume without loss of generality that $P \geq 1/2$. If we define the probability in absentia of noise to be P_0 and the probability of an error occurring to be $P_e < 1$ then

$$P = P_0(1 - P_e) + P_e/2. \tag{F1}$$

Now let b be the value of a bit inferred from such an experiment and let n be the number of shots used to infer this value by majority voting. The Chernoff bound then reads

$$P(b = 0|P_0 > 1/2) \leq e^{\frac{-n}{2P}(P-1/2)^2}. \tag{F2}$$

Note that the assumptions that underlie the Chernoff bound are met because $P_0 > 1/2$ and $P_e < 1$ because P is a monotonic function of P_0. As P_e tends to 1, P approaches $1/2$. This causes the exponent in the Chernoff bound to tend to 0. Thus we no longer gain exponential suppression from n in this limit; whereas, for small P_e the exponent is approximately $\frac{1}{2}\frac{(P_0-1/2)^2}{P_0} - O(P_e)$. Thus the exponential suppression of errors expected from majority voting will hold so long as P_e is small and P_0 is bounded away from $1/2$.

Now if we assume that the number of bits needed is n_{bits} then the mean number of bits that are incorrect is $n_{\text{bits}}e^{\frac{-n}{2P}(P-1/2)^2}$, assuming $P_0 > 1/2$ for each bit being inferred. If we solve for the case where the mean number of

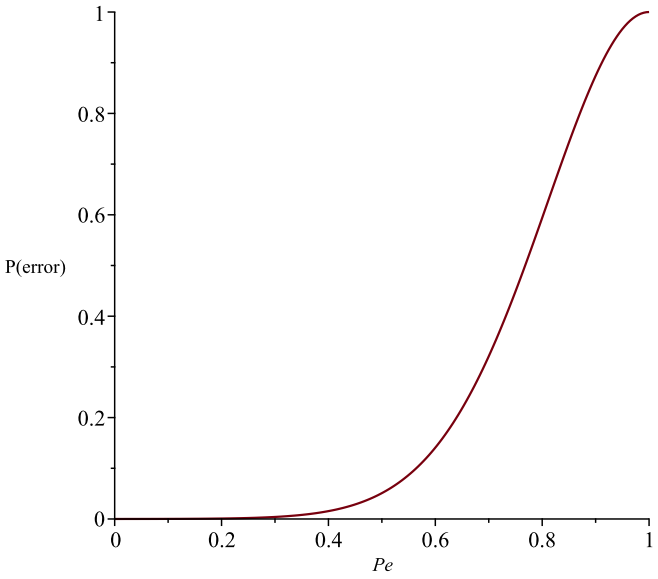

FIG. 9. Upper bound on error probability of inferring a single bit in a phase estimation experiment given 500 shots and $P_0 = 2/3$. The data shows substantial suppression of errors until $P_e \approx 1/3$ and then rapidly diverges thereafter.

erroneous bits is at most 1 we find that

$$P_0 \geq \frac{1}{2} + \frac{\sqrt{n_{\text{bits}} \ln(n_{\text{bits}}) + \ln^2(n_{\text{bits}})} - \ln(n_{\text{bits}})}{n(P_e - 1)}, \tag{F3}$$

suffices to guarantee that this condition is met. This shows that the critical signal that we have to have in order to have a reasonable probability of inferring all of the bits rapidly diverges as P_e approaches 1. Furthermore, as the number of bits inferred increases the number of shots needed for each inferred bit must increase; although this increase is modest with respect to the number of bits required.

To visualize how this affect plays out, we consider a plot of the probability of erring in the inference of a bit using a sample of 500 shots (which is easily achievable in our apparatus). We further assume a value of $P_0 = 2/3$, which again is roughly typical of phase estimation experiments. We observe that this simple analysis from the Chernoff bound qualitatively reproduces the curves observed in the main body for the error as a function of T_2 or the heater variance. This suggests that the rapid breakdown of IPEA can be understood as a simple consequence of majority voting and the inability of IPEA to detect inference errors and adapt after such an error occurs.

Appendix G: Rescaling the decoherence time T_2

In the decoherence model used T_2 represents the decoherence time in units of time required to perform each of the M controlled gates that would be needed to obtain the controlled-U^M operation in a scalable manner. Its value in standard units thus depends on the particular architecture exploited. For example, in a superconductive qubits architecture the operational time for a three qubit controlled-U (using a single Trotter step) is typically $\simeq 1.5\ \mu s$ [9]. From our study, this implies that a fully scalable 16-bits IPEA implementation starts to be impractical in superconductive devices with decoherence time approximately lower than 50 μs. Given that the relaxation time of state-of-the-art superconductive qubits is on the same order of magnitude, this highlights how the implementation of IPEA on current pre-fault tolerant devices is very limited. An equivalent implementation of IPEA on a solid-state electron-spin system would require approximately 50 μs for a single controlled-U [28], and thus 1.5 ms total time for a fully scalable 16-bits IPEA. For such systems, typically observed decoherence times T_2 are of the order of tens of ms.